

THE F5 CRITERION REVISED

ALBERTO ARRI AND JOHN PERRY

ABSTRACT. The purpose of this work is to generalize part of the theory behind Faugère’s “F5” algorithm. This is one of the fastest known algorithms to compute a Gröbner basis of a polynomial ideal I generated by polynomials f_1, \dots, f_m . A major reason for this is what Faugère called the algorithm’s “new” criterion, and we call “the F5 criterion”; it provides a sufficient condition for a set of polynomials G to be a Gröbner basis. However, the F5 algorithm is difficult to grasp, and there are unresolved questions regarding its termination.

This paper introduces some new concepts that place the criterion in a more general setting: \mathcal{S} -Gröbner bases and primitive \mathcal{S} -irreducible polynomials. We use these to propose a new, simple algorithm based on a revised F5 criterion. The new concepts also enable us to remove various restrictions, such as proving termination without the requirement that f_1, \dots, f_m be a regular sequence.

1. INTRODUCTION

Since their introduction in [3] by B. Buchberger, Gröbner bases and their computation have attracted significant attention in the computer algebra community. The best-known algorithm used to compute a Gröbner basis is the original algorithm due to Buchberger [3], and named after him. Its efficiency has been constantly enhanced through the years, but there remains room for improvement. Various criteria have since been introduced to detect useless computations – for example, [3, 4, 9] — but even so, the algorithm spends most of its time reducing polynomials to zero (“zero reductions”).

In [11], Lazard pointed out that one can view the computation of a Gröbner basis as the reduction to row-echelon form of the Macaulay matrix of the ideal. This led to the Staggered Linear Basis algorithm of Gebauer and Möller in [8], as well as the “F4” algorithm of Faugère in [6]. Möller, Mora, and Traverso exploited the relationship between zero reductions and syzygies [12], but although the algorithm they presented successfully detected many zero reductions, in practice it took too much memory and time (see Section 8 of [12]). In [5], Faugère combined aspects of these approaches into algorithm “F5”, which for a certain class of polynomial system eliminates *all* zero reductions. This algorithm exhibits impressive performance.

By Faugère’s admission, the theory behind the algorithm’s new criterion, which we call *the F5 criterion*, is merely sketched, so as to leave more room for examples and an accurate description of the algorithm. The proof of the algorithm’s termination and correctness were likewise only outlined. Additionally, some arguments were made under strong assumptions, such as that the input sequence f_1, \dots, f_m had only principal syzygies (such a sequence is called a *regular* sequence).

We pause a moment to consider some variants of F5. Bardet described an implementation of F5 in matrix form, where termination is ensured by manually supplying a maximal degree [2]. Stegers filled in some details of Faugère’s proof in [15], but stopped at two conjectures, one of which Gash later showed to be false [7].

The purpose of this paper is to present a simpler algorithm that illustrates the fundamental principles of F5 without sacrificing termination. We begin by defining a function \mathcal{S} which is equivalent to that of Faugère, then develop a structured theory, introducing new concepts such as *primitive \mathcal{S} -irreducible polynomials* and *\mathcal{S} -Gröbner bases*. These make the study of the problem more accessible, and suggest a new version of the F5 criterion which depends neither on the regularity of the input, nor on a particular ordering on the module of syzygies.

From this theory, we develop a new, simpler algorithm. We must emphasize that the algorithm is a simple demonstration of the criterion, and not a deep treatment of how to implement a highly efficient algorithm; nevertheless, the new concepts allow us to prove correctness and termination *for any input*. Note that although some F5-style algorithms provide explicit termination mechanisms [2, 7], these mechanisms rely on previously-developed, non-F5 criteria to compute explicitly a maximal degree; by contrast, the termination criterion used here is precisely the generalized F5 criterion used to detect useless computations. Later, we show that if we know that the input is a regular sequence and we use a specific ordering on $\text{Syz } \mathcal{F}$, we can avoid all the reductions to zero. We compare the results to both F5 and the Staggered Linear Basis algorithm, showing how this new algorithm differs from each.

The paper's structure is as follows. Sections 2–4 cover background material; although most of this is relatively straightforward, an important and novel contribution of the paper appears at the end of Section 4 with Proposition 14. The proof of that theorem leads to the concept of *primitive \mathcal{S} -irreducible polynomials*, from which we obtain in Section 5 a new characterization theorem for a Gröbner basis (Theorem 18). In Section 6, we use this characterization to formulate the new algorithm, and we prove that it terminates correctly. Section 7 compares this algorithm to the Staggered Linear Basis algorithm and F5, illustrating the differences concretely. Section 8 describes some conclusions and possible future directions.

2. PRELIMINARIES

Let $P = k[x_1, \dots, x_n]$ be the polynomial ring over the field k with n indeterminates, let μ be any admissible ordering on \mathbb{T}^n , the monoid of power products over x_1, \dots, x_n : $\mathbb{T}^n = \{\prod_{i=1}^n x_i^{\alpha_i} \mid \alpha_i \in \mathbb{N}\}$.

Let P^m be the free P -module generated by $\{e_1, \dots, e_m\}$ and let μ' be any admissible¹ ordering on \mathbb{T}_m^n , the set of module terms of P^m : $\mathbb{T}_m^n = \{te_l \mid t \in \mathbb{T}, l \in \{1, \dots, m\}\}$.

Fix $\mathcal{F} = (f_1, \dots, f_m) \in P^m$ and let $I \subseteq P$ be the ideal generated by \mathcal{F} , and define $v : P^m \rightarrow I$ as the P -module homomorphism such that $v(e_i) = f_i$, and let $\text{Syz } \mathcal{F} = \ker v$, so that $\text{Syz } \mathcal{F}$ is the module of syzygies of \mathcal{F} , $\text{LT}(\text{Syz } \mathcal{F}) \subseteq \mathbb{T}_m^n$ is set of leading module terms of $\text{Syz } \mathcal{F}$, and $\text{NS}(\text{Syz } \mathcal{F}) = \mathbb{T}_m^n \setminus \text{LT}(\text{Syz } \mathcal{F})$ is the *normal set* of the syzygies of \mathcal{F} .

Clearly v is surjective; therefore, as a P -module, $P^m/\text{Syz } \mathcal{F} \simeq I$. Let $\psi : I \rightarrow P^m/\text{Syz } \mathcal{F}$ be the P -module isomorphism between them. We use the notation $\text{LT}(\cdot)$ for both the leading term of a polynomial in P with respect to μ , and the module leading term of a module element in P^m with respect to μ' . We will use $\text{LC}(f)$, where f is a nonzero polynomial belonging to I , to denote the coefficient of $\text{LT}(f)$.

We are interested in finding a set of polynomials G such that G is a Gröbner basis for I with respect to the ordering μ on \mathbb{T}^n .

¹On what we mean by an “admissible” ordering, see the appendix.

Definition 1. Let

$$\begin{aligned} \mathcal{S} : I \setminus \{0\} &\rightarrow \text{NS}(\text{Syz } \mathcal{F}) \\ f &\mapsto \text{LT}(\psi(f)), \end{aligned}$$

where $\text{LT}(\psi(f))$ is the module leading term of the normal form of $\psi(f)$ with respect to the ordering μ' on P^m .

The key idea of Faugère is to keep track of the value of $\mathcal{S}(f)$ for any polynomial f we will work with. It is however clear from the definition that the explicit calculation of \mathcal{S} requires, at least, to know a Gröbner basis of $\text{Syz } \mathcal{F}$ which is computationally expensive to compute, more than a Gröbner basis of I itself. In fact, we will obtain \mathcal{S} from the fact that $\mathcal{S}(f_i) = e_i$ (unless $e_i \in \text{LT}(\text{Syz } \mathcal{F})$) and from other properties of \mathcal{S} .

3. PROPERTIES OF \mathcal{S}

Lemma 2 (Properties of \mathcal{S}). *Let $f, f_1, f_2 \in I \setminus \{0\}$. The following hold:*

(1) *If $\mathcal{S}(f_1) > \mathcal{S}(f_2)$ then:*

$$\mathcal{S}(f_1 + f_2) = \mathcal{S}(f_1).$$

(2) *If $\mathcal{S}(f_1) = \mathcal{S}(f_2) = \sigma$ and there is no $\lambda \in k^* = k \setminus \{0\}$ such that $f_1 = \lambda f_2$, then there exist α and β in k^* such that:*

$$\mathcal{S}(\alpha f_1 + \beta f_2) < \sigma.$$

(3) *Let $t \in \mathbb{T}^n$, then*

$$\begin{aligned} \mathcal{S}(tf) = t\mathcal{S}(f) &\iff t\mathcal{S}(f) \in \text{NS}(\text{Syz } \mathcal{F}), \\ \mathcal{S}(tf) < t\mathcal{S}(f) &\iff t\mathcal{S}(f) \in \text{LT}(\text{Syz } \mathcal{F}). \end{aligned}$$

Proof. (1) and (2) are trivial.

In order to prove (3), let $\mathcal{S}(f) = \sigma = \tau e_i$. By the definition of \mathcal{S} we have:

$$f = v(\alpha \tau e_i + \text{smaller terms}),$$

where $\alpha \in k^*$, $\tau \in \mathbb{T}^n$, and the argument of v is in its normal form with respect to $\text{Syz } \mathcal{F}$.

Multiplying both sides by t , we get:

$$tf = v(\alpha t \tau e_i + \text{smaller terms}).$$

If $t \tau e_i = t \sigma \in \text{NS}(\text{Syz } \mathcal{F})$, the leading term of the normal form of $\alpha t \tau e_i + \dots$ is $t \sigma$ and, in this case, $\mathcal{S}(tf) = t \sigma = t \mathcal{S}(f)$. Otherwise, $t \sigma \notin \text{NS}(\text{Syz } \mathcal{F})$, so the normal form has a leading term which is strictly smaller than $t \sigma$ and we have $\mathcal{S}(tf) < t \mathcal{S}(f)$. \square

Corollary 3. *To decide whether $\mathcal{S}(tf) = t\mathcal{S}(f)$, it suffices to know $\text{NS}(\text{Syz } \mathcal{F})$ or, equivalently, $\text{LT}(\text{Syz } \mathcal{F})$. Also, if $t\mathcal{S}(f) = \mathcal{S}(g)$ for some $g \in I \setminus \{0\}$, then since $\mathcal{S}(g) \in \text{NS}(\text{Syz } \mathcal{F})$, we can conclude that $t\mathcal{S}(f) = \mathcal{S}(tf)$.*

One of the key concepts of the classic theory of Gröbner bases is the polynomial reduction: one says that $f \in P \setminus \{0\}$ reduces with a $h \in P \setminus \{0\}$, if there exist $\alpha \in k^*$ and $t \in \mathbb{T}^n$ such that $\text{LT}(f - \alpha th) < \text{LT}(f)$, denoted

$$f \xrightarrow{h} g$$

where $g = f - \alpha th$.

We now introduce a special kind of reduction for a polynomial f , which takes in consideration the value of $\mathcal{S}(f)$.

Definition 4 (\mathcal{S} -reduction). Let $f, h \in I \setminus \{0\}$, $g \in I$ and $\sigma \in \mathbb{T}_m^n$. We say that f \mathcal{S} -reduces with respect to σ to g with h ,

$$f \xrightarrow{h}_{\mathcal{S}, \sigma} g$$

if there are $t \in \mathbb{T}^n$ and $\alpha \in k^*$ such that:

- $\text{LT}(g) < \text{LT}(f)$ and $f - \alpha th = g$, and
- $\mathcal{S}(th) < \sigma$.

When we omit to specify σ , we assume $\sigma = \mathcal{S}(f)$.

Note that this reduction is defined only for polynomials f which belong to the ideal I , and not for arbitrary elements of the ring P . Also, when $\sigma = \mathcal{S}(f)$, since $\mathcal{S}(th) < \mathcal{S}(f)$ we have $\mathcal{S}(g) = \mathcal{S}(f - \alpha th) = \mathcal{S}(f)$. Hence, when performing one, or more, \mathcal{S} -reduction steps with a polynomial:

$$f \xrightarrow{h_0}_{\mathcal{S}} f_1 \xrightarrow{h_1}_{\mathcal{S}} f_2 \xrightarrow{h_2}_{\mathcal{S}} \dots \xrightarrow{h_{k-1}}_{\mathcal{S}} f_k$$

we have $\mathcal{S}(f) = \mathcal{S}(f_1) = \dots = \mathcal{S}(f_k)$ and $\text{LT}(f) > \text{LT}(f_1) > \dots > \text{LT}(f_k)$; that is, the value of \mathcal{S} is kept constant, while the leading term decreases.

Let us consider how to characterize those elements which cannot be further \mathcal{S} -reduced with respect to a given $\sigma \in \mathbb{T}_m^n$. The following definition is natural:

Definition 5 (\mathcal{S} -irreducible polynomial). We say that $f \in I$ is \mathcal{S} -irreducible with respect to $\sigma \in \mathbb{T}_m^n$ if $f = 0$ or if there is no $h \in I$ which \mathcal{S} -reduces f with respect to σ . As before, if we do not specify σ , we assume $\sigma = \mathcal{S}(f)$. Note that this definition depends on the values of I , \mathcal{F} and μ' .

We could look for a criterion which decides whether a given set of nonzero polynomials G is a Gröbner basis by looking at the values of $\mathcal{S}(g)$ for all g in G . However, it is wiser to characterize a set of polynomials with a property similar to that of a Gröbner basis, but which also accounts for \mathcal{S} . We therefore introduce the following:

Definition 6 (\mathcal{S} -Gröbner basis). We say that $G \subset I$ is an \mathcal{S} -Gröbner basis if for each \mathcal{S} -irreducible polynomial $f \in I \setminus \{0\}$, there exist $g \in G$ and $t \in \mathbb{T}^n$ such that $\text{LT}(tg) = \text{LT}(f)$ and $\mathcal{S}(tg) = \mathcal{S}(f)$.

Remark 7. An \mathcal{S} -Gröbner basis depends on:

- the ideal I ,
- the term ordering μ on \mathbb{T}^n ,
- the m -tuple of generators \mathcal{F} ,
- the ordering μ' on \mathbb{T}_m^n .

We will prove in the following section that an \mathcal{S} -Gröbner basis is a Gröbner basis in the usual sense. While Definition 6 is not especially useful from a computational point of view, inasmuch as it is quantified over an infinite set, Theorem 18 will provide us an equivalent criterion that is quantified over a finite set. Before we can prove it, however, we need to consider some properties of \mathcal{S} -reductions.

4. PROPERTIES OF \mathcal{S} -REDUCTIONS

In this section we will prove the main facts which will lead to the characterization we are looking for.

Definition 8. Let

$$\begin{aligned} \varphi : \text{LT}(I) &\rightarrow \text{NS}(\text{Syz } \mathcal{F}) \\ t &\mapsto \min\{\mathcal{S}(f) \mid f \in I, \text{LT}(f) = t\}. \end{aligned}$$

In other words, if t belongs to $\text{LT}(I)$, φ is the minimum value \mathcal{S} can take on a polynomial whose leading term is t . It follows that, for any $f \in I \setminus \{0\}$, $\varphi(\text{LT}(f)) \leq \mathcal{S}(f)$ always holds.

Lemma 9. φ is a bijection, and the inverse function of φ has an explicit formula: $\varphi^{-1}(\sigma) = \min\{t' \in \mathbb{T}^n \mid \exists f \in I, \text{LT}(f) = t', \mathcal{S}(f) = \sigma\}$.

Proof. We show that φ is both injective and surjective.

Injective:: By way of contradiction, suppose there exist $\sigma \in \text{NS}(\text{Syz } \mathcal{F})$ and $t_1, t_2 \in \text{LT}(I)$ such that $t_1 > t_2$ and $\sigma = \varphi(t_1) = \varphi(t_2)$. Then we can find $f_1, f_2 \in I$ such that $\text{LT}(f_1) = t_1$, $\text{LT}(f_2) = t_2$, and $\mathcal{S}(f_1) = \mathcal{S}(f_2) = \sigma$. By Lemma 2, there exist $\alpha, \beta \in k^*$ such that $\mathcal{S}(\alpha f_1 + \beta f_2) < \sigma$, but $\text{LT}(\alpha f_1 + \beta f_2) = t_1$, and therefore $\varphi(t_1) < \sigma$, contradicting the hypothesis.

Surjective:: Let $\sigma = \tau e_i \in \text{NS}(\text{Syz } \mathcal{F})$, define

$$t = \min\{t' \in \mathbb{T}^n \mid \exists f \in I, \text{LT}(f) = t', \mathcal{S}(f) = \sigma\}.$$

(This set is not empty because it contains $\text{LT}(\tau f_i)$.) Let $f \in I$ be a polynomial with $\text{LT}(f) = t$ and $\mathcal{S}(f) = \sigma$; obviously $\varphi(t) \leq \sigma$. By way of contradiction, suppose that $\varphi(t) < \sigma$. Then there exists $f' \in I$ such that $\text{LT}(f') = t$ and $\mathcal{S}(f') < \sigma$. We can now choose $\alpha, \alpha' \in k^*$ with $\text{LT}(\alpha f + \alpha' f') < t$ such that $\mathcal{S}(\alpha f + \alpha' f') = \sigma$. The existence of $\alpha f + \alpha' f'$ contradicts the minimality of t ; therefore, $\varphi(t) = \sigma$. □

The fact that φ is a bijection will play a crucial role in most of the subsequent proofs.

Theorem 10 (\mathcal{S} -reduction theorem). *Let $f \in I$ and $\sigma \in \mathbb{T}_m^n$ such that f is \mathcal{S} -irreducible with respect to σ and f is of the form $f = v(\alpha\sigma + \text{smaller terms})$, for some $\alpha \in k^*$.*

Either the following equivalent propositions hold:

- (a) $f = 0$,
- (b) $\sigma \in \text{LT}(\text{Syz } \mathcal{F})$,

or the following equivalent propositions hold:

- (1) $f \neq 0$,
- (2) $\sigma \in \text{NS}(\text{Syz } \mathcal{F})$,
- (3) $f \neq 0$ and $\sigma = \mathcal{S}(f) = \varphi(\text{LT}(f))$.

Proof.

a \Rightarrow b: Suppose $f = 0$, then $0 = f = v(\alpha\sigma + \dots)$. It follows that $\sigma \in \text{LT}(\text{Syz } \mathcal{F})$.

b \Rightarrow a: Assume by way of contradiction that $\sigma \in \text{LT}(\text{Syz } \mathcal{F})$ and $f \neq 0$. Let $t = \text{LT}(f)$, and consider $\sigma' = \varphi(t) \in \text{NS}(\text{Syz } \mathcal{F})$. There exists $g \in I$ such that $\text{LT}(g) = t$ and $\mathcal{S}(g) = \sigma'$; since $\sigma' < \sigma$, g is an \mathcal{S} -reducer for f with respect to σ , contradicting the fact that f is \mathcal{S} -irreducible. Therefore, $f = 0$.

1 \Rightarrow 2: Assume by way of contradiction that $f \neq 0$ and $\sigma \notin \text{NS}(\text{Syz } \mathcal{F})$. Then $\sigma \in \text{LT}(\text{Syz } \mathcal{F})$. Let $t = \text{LT}(f)$, and consider $\sigma' = \varphi(t) \in \text{NS}(\text{Syz } \mathcal{F})$. There exists $g \in I$ such that $\text{LT}(g) = t$ and $\mathcal{S}(g) = \sigma'$; since $\sigma' < \sigma$, g is an \mathcal{S} -reducer for f with respect to σ , contradicting the hypothesis that f is \mathcal{S} -irreducible. Therefore, $\sigma \in \text{NS}(\text{Syz } \mathcal{F})$.

2 \Rightarrow 3: Assume $\sigma \in \text{NS}(\text{Syz } \mathcal{F})$. Necessarily, $\sigma = \mathcal{S}(f)$. Suppose now that $\sigma \neq \varphi(\text{LT}(f))$; then $\mathcal{S}(f) > \varphi(\text{LT}(f))$. Therefore there exists a polynomial $g \in I$ such that $t = \text{LT}(g) = \text{LT}(f)$ and $\varphi(t) = \mathcal{S}(g) = \varphi(\text{LT}(f)) < \mathcal{S}(f)$. It follows that g is an \mathcal{S} -reductor of f , and f is not \mathcal{S} -irreducible.

3 \Rightarrow 1: Obvious. □

Theorem 10 implies that it only makes sense to consider those polynomials f that are \mathcal{S} -irreducible with respect to $\mathcal{S}(f)$. Also, an \mathcal{S} -reduction yields 0 if and only if performed with respect to a $\sigma \in \text{LT}(\text{Syz } \mathcal{F})$; conversely, if an \mathcal{S} -reduction yields a non-zero polynomial, then we know that it was performed with respect to some $\sigma \in \text{NS}(\text{Syz } \mathcal{F})$.

Remark 11. Observe that a polynomial f is \mathcal{S} -irreducible iff $\mathcal{S}(f) = \varphi(\text{LT}(f))$; otherwise, $\mathcal{S}(f) > \varphi(\text{LT}(f))$, and we could find $g \in I$ such that $\text{LT}(g) = \text{LT}(f)$ and $\mathcal{S}(g) = \varphi(\text{LT}(f))$, so that g would \mathcal{S} -reduce f .

In strict analogy with the classic Gröbner basis theory we have the following result:

Proposition 12. *If G is an \mathcal{S} -Gröbner basis then for any nonzero $f \in I$ such that f is not \mathcal{S} -irreducible, there exists $g \in G$ and $t \in \mathbb{T}^n$ such that:*

- $\text{LT}(tg) = \text{LT}(f)$,
- $\mathcal{S}(tg) = t\mathcal{S}(g) < \mathcal{S}(f)$.

That is, it is always possible to find an \mathcal{S} -reductor for f in G .

Proof. Since f is not \mathcal{S} -irreducible, take h \mathcal{S} -irreducible such that $\text{LT}(f) = \text{LT}(h)$. From the remark above, $\mathcal{S}(h) < \mathcal{S}(f)$, so h is an \mathcal{S} -reductor of f . We can then find $t \in \mathbb{T}$ and $g \in G$ such that $t \text{LT}(g) = \text{LT}(h) = \text{LT}(f)$ and (using Corollary 3) $\mathcal{S}(tg) = t\mathcal{S}(g) = \mathcal{S}(h)$. □

This fact combined with lemma 9 leads immediately to:

Proposition 13. *If G is an \mathcal{S} -Gröbner basis, then G is a Gröbner basis with respect to the ordering μ on \mathbb{T}^n .*

Proof. For any $t \in \text{LT}(I)$, Lemma 9 implies that there exists $\sigma \in \mathbb{T}_m^n$ such that $\varphi^{-1}(\sigma) = t$. Let $f \in I$ such that $\text{LT}(f) = t$ and $\mathcal{S}(f) = \sigma$. From Proposition 12, we may assume that f is \mathcal{S} -irreducible (if not, \mathcal{S} -reduce it). Then $\exists g \in G, u \in \mathbb{T}^n$ such that $\text{LT}(ug) = \text{LT}(f) = t$. Hence the set $\{\text{LT}(g) \mid g \in G\}$ generates $\text{LT}(I)$ and $(G) \subseteq I$. Therefore G is a Gröbner basis for I . □

Proposition 14. *Every \mathcal{S} -Gröbner basis contains a finite \mathcal{S} -Gröbner basis.*

Proof. Let $G = \{g_i\}_{i \in \mathcal{I}}$ be an \mathcal{S} -Gröbner basis, and

$$\begin{aligned} \vartheta : G &\rightarrow P \oplus P^m \\ g_i &\mapsto (\text{LT}(g_i), \mathcal{S}(g_i)). \end{aligned}$$

The image $\vartheta(G)$ generates a P -submodule M of the P -module $P \oplus (P)^m \cong (P)^{m+1}$. This is a noetherian module; therefore, there exists a finite subset \mathcal{J} of \mathcal{I} such that $\vartheta(G')$ generates M , for some $G' = \{g_j\}_{j \in \mathcal{J}}$.

We claim that G' is itself an \mathcal{S} -Gröbner basis. To see this, let $f \in I$ be an \mathcal{S} -irreducible polynomial. By definition, $\mathcal{S}(f) \in \text{NS}(\text{Syz } \mathcal{F})$. Since G is an \mathcal{S} -Gröbner basis, we can find a $g_i \in G$ and a $t \in \mathbb{T}^n$ such that $t\mathcal{S}(g_i) = \mathcal{S}(tg_i) = \mathcal{S}(f)$ and $t \text{LT}(g_i) = \text{LT}(tg_i) = \text{LT}(f)$

(using Lemma 2(3) for $t\mathcal{S}(g_i) = \mathcal{S}(tg_i)$). If $i \in \mathcal{J}$, then $g_i \in G'$ and we're fine. Otherwise, $i \in \mathcal{I} \setminus \mathcal{J}$; since $\vartheta(g_i) \in M$, there exist $j_i \in \mathcal{J}$ and $t_i \in \mathbb{T}^n$ such that

$$t_i\vartheta(g_{j_i}) = \vartheta(g_i).$$

This implies that $\mathcal{S}(g_i) = t_i\mathcal{S}(g_{j_i}) = \mathcal{S}(t_i g_{j_i})$ and $\text{LT}(g_i) = \text{LT}(t_i g_{j_i})$. Then $\mathcal{S}(tg_i) = \mathcal{S}(t(t_i g_{j_i}))$ and $\text{LT}(tg_i) = \text{LT}(t(t_i g_{j_i}))$, so we have found $g_{j_i} \in G'$ and $t \cdot t_i \in \mathbb{T}^n$ which satisfy the \mathcal{S} -Gröbner basis property for f . \square

The elements of G' from the proof above will prove critically important when we examine our algorithm, so we will identify them by a special term.

Definition 15 (Primitive \mathcal{S} -irreducible polynomial). We say that a nonzero polynomial f \mathcal{S} -irreducible with respect to $\mathcal{S}(f)$ is *primitive \mathcal{S} -irreducible* if there are no polynomials $f' \in I \setminus \{0\}$ and terms $t \in \mathbb{T}^n \setminus \{1\}$ such that f' is \mathcal{S} -irreducible, $\text{LT}(tf') = \text{LT}(f)$ and $\mathcal{S}(tf') = \mathcal{S}(f)$.

The proof of Proposition 14 implies that if we have an \mathcal{S} -Gröbner basis G , then we can obtain a finite \mathcal{S} -Gröbner basis by keeping a subset of primitive \mathcal{S} -irreducible polynomials with different leading terms. Hence there exist \mathcal{S} -Gröbner bases which contain only primitive \mathcal{S} -irreducible polynomials.

5. THE MAIN RESULT

First we adapt the definition of a normal pair in [5] to reflect primitive \mathcal{S} -irreducible polynomials.

Definition 16 (Normal Pair). Given $g_1, g_2 \in I \setminus \{0\}$, let $\text{Spol}(g_1, g_2) = u_1 g_1 - u_2 g_2$ be the S-polynomial of g_1 and g_2 ; that is, $u_i = \frac{\text{lcm}(\text{LT}(g_1), \text{LT}(g_2))}{\text{LC}(g_i) \text{LT}(g_i)}$. We say that (g_1, g_2) is a *normal pair* if:

- (1) g_i is a primitive \mathcal{S} -irreducible polynomial for $i = 1, 2$,
- (2) $\mathcal{S}(u_i g_i) = \text{LT}(u_i) \mathcal{S}(g_i)$ for $i = 1, 2$,
- (3) $\mathcal{S}(u_1 g_1) \neq \mathcal{S}(u_2 g_2)$.

Remark 17. With this definition, if (g_1, g_2) is a normal pair, then

$$\mathcal{S}(\text{Spol}(g_1, g_2)) = \max(\mathcal{S}(u_1 g_1), \mathcal{S}(u_2 g_2))$$

will always hold. In addition, if $\mathcal{S}(u_1 g_1) > \mathcal{S}(u_2 g_2)$, then $u_1 \neq 1$, as if u_1 were 1, g_2 would be an \mathcal{S} -reductor of g_1 . Therefore $\mathcal{S}(\text{Spol}(g_1, g_2)) > \max(\mathcal{S}(g_1), \mathcal{S}(g_2))$.

Theorem 18 (F5 criterion). *Suppose that G is a set of \mathcal{S} -irreducible polynomials of I , such that:*

- *for each $i = 1, \dots, m$ such that $e_i \notin \text{LT}(\text{Syz } \mathcal{F})$ there exists $g_i \in G$ such that $\mathcal{S}(g_i) = e_i$, and*
- *for any $g_1, g_2 \in G$ such that (g_1, g_2) is a normal pair, there exist $g \in G$ and $t \in \mathbb{T}^n$ such that tg is \mathcal{S} -irreducible and $\mathcal{S}(tg) = \mathcal{S}(\text{Spol}(g_1, g_2))$.*

Then G is a \mathcal{S} -Gröbner basis of I .

Remark 19 (Rewritable criterion). Note that the second condition does not explicitly involve the S-polynomial of a pair (g_1, g_2) , but cares only about $\mathcal{S}(\text{Spol}(g_1, g_2))$. Hence, we can think of this as a criterion to choose elements of $\text{NS}(\text{Syz } \mathcal{F})$ instead of polynomials. Additionally,

if two or more normal pairs are such that \mathcal{S} takes the same value on their S-polynomials, we can freely consider just one of them.

Proof. As noted at the end of the previous section, we may, without loss of generality, assume that the elements of G are primitive \mathcal{S} -irreducible and have distinct leading terms. By way of contradiction, suppose that there exists a minimal $\sigma \in \text{NS}(\text{Syz } \mathcal{F})$ and an \mathcal{S} -irreducible $f \in I \setminus \{0\}$ with $\mathcal{S}(f) = \sigma$ and the \mathcal{S} -Gröbner basis property does not hold for f and σ . That is, for all $g \in G$ and for all $t \in \mathbb{T}^n$, $\text{LT}(tg) \neq \text{LT}(f)$ or $\mathcal{S}(tg) \neq \mathcal{S}(f)$.

The first hypothesis implies that there exist at least one primitive \mathcal{S} -irreducible $g \in G$ and some $\tau \in \mathbb{T}^n$ such that $\tau\mathcal{S}(g) = \mathcal{S}(f) = \sigma$; among the possible choices for g and τ , pick one which minimizes $\text{LT}(\tau g)$. By Lemma 2(3), $\mathcal{S}(\tau g) = \tau\mathcal{S}(g) = \sigma$. Hence $\text{LT}(\tau g) \neq \text{LT}(f)$. By Remark 11, $\mathcal{S}(f) = \varphi(\text{LT}(f))$, and by Lemma 9, $\text{LT}(\tau g) > \text{LT}(f)$. In addition, we have $\mathcal{S}(\tau g) = \mathcal{S}(f) = \varphi(\text{LT}(f)) \neq \varphi(\text{LT}(\tau g))$, so again by Remark 11, τg is not \mathcal{S} -irreducible.

By Lemma 2(2), there exist $\alpha, \beta \in k^*$ such that $\mathcal{S}(\alpha f + \beta \tau g) = \sigma'$ for some $\sigma' < \sigma$. Since σ was chosen to be the minimal element of $\text{NS}(\text{Syz } \mathcal{F})$ such that the \mathcal{S} -Gröbner basis property does not hold, Definition 6 and Proposition 12 applied to $\alpha f + \beta \tau g$ imply that there exist $g' \in G$ and $\tau' \in \mathbb{T}^n$ such that $\text{LT}(\tau' g') = \text{LT}(\alpha f + \beta \tau g) = \text{LT}(\tau g)$ and

$$\mathcal{S}(\tau' g') = \tau' \mathcal{S}(g') \leq \mathcal{S}(\alpha f + \beta \tau g) = \sigma' < \sigma = \mathcal{S}(\tau g).$$

Clearly $g \neq g'$.

It follows that (g, g') is a normal pair. From the second hypothesis, we know that there exist $g'' \in G$ and $\tau'' \in \mathbb{T}^n$ such that $\tau'' g''$ is \mathcal{S} -irreducible and $\mathcal{S}(\tau'' g'') = \mathcal{S}(\text{Spol}(g, g'))$. Write $\hat{\tau} \text{Spol}(g, g') = \gamma \tau g - \gamma' \tau' g'$, for some $\gamma, \gamma' \in k^*$, where $\hat{\tau}$ is the gcd of τ and τ' . Since (g, g') is a normal pair and $\sigma \in \text{NS}(\text{Syz } \mathcal{F})$,

$$\sigma = \tau \mathcal{S}(g) = \hat{\tau} \mathcal{S}(\text{Spol}(g, g')) = \hat{\tau} \mathcal{S}(\tau'' g'') = \mathcal{S}(\hat{\tau} \tau'' g'').$$

By Remark 11, $\mathcal{S}(\tau'' g'') = \varphi(\text{LT}(\tau'' g''))$, so we have

$$\text{LT}(\tau'' g'') = \varphi^{-1}(\mathcal{S}(\tau'' g'')) \leq \text{LT}(\text{Spol}(g, g')).$$

Multiplying both sides by $\hat{\tau}$, we have

$$\text{LT}(\hat{\tau} \tau'' g'') \leq \text{LT}(\gamma \tau g - \gamma' \tau' g') < \text{LT}(\tau g).$$

The existence of g'' and $\hat{\tau} \tau''$ contradicts the choice of g and τ . □

6. THE ALGORITHM

We shall now present a simple algorithm which computes as \mathcal{S} -Gröbner basis of an ideal based on the criterion. This algorithm is quite different from Faugère's, in that it is a direct application of the criterion. In particular, it does not involve reductions that yield more than one result, nor the more rigorous simplification rules. See Section 7.2 for a detailed discussion.

One first problem is that to check condition 2 of definition 16 we need to know $\text{LT}(\text{Syz } \mathcal{F})$, since

$$\mathcal{S}(tf) = t\mathcal{S}(f) \iff t\mathcal{S}(f) \notin \text{LT}(\text{Syz } \mathcal{F}).$$

We almost never know this before hand; therefore, we introduce a new variable L , a subset of $\text{LT}(\text{Syz } \mathcal{F})$. At the beginning of the algorithm, we simply assume $L = \emptyset$. We make use

of L whenever we need to check if $t\mathcal{S}(f) = \mathcal{S}(tf)$ by checking whether $t\mathcal{S}(f)$ belongs to $\langle L \rangle \subseteq P^m$, the P -module generated by L . We then replace condition 2 of definition 16 by:

$$\mathcal{S}(u_i g_i) = \text{LT}(u_i) \mathcal{S}(g_i) \iff \text{LT}(u_i) \mathcal{S}(g_i) \notin \langle L \rangle.$$

By doing so, we end up considering more pairs than we should, but we do not skip any legitimate pair.

So, when (g_1, g_2) is a normal pair (with the weakened condition 2), we calculate a polynomial $f = \text{Spol}(g_1, g_2)$ and a $\sigma = \max(u_1 g_1, u_2 g_2)$. Thereafter we \mathcal{S} -reduce f with respect to σ . Note that σ satisfies the hypothesis of Theorem 10. If the \mathcal{S} -reduction yields 0, we know that $\sigma \in \text{LT}(\text{Syz } \mathcal{F})$; accordingly, we enlarge L by inserting σ . Otherwise, we obtain a nonzero polynomial, which tells us that $\sigma = \mathcal{S}(f)$.

G is the set which will contain the \mathcal{S} -Gröbner basis; we add elements to G as we find them. For each element g we add to G , we also store $\mathcal{S}(g)$; thus, G is more precisely a set of pairs (g, σ) . When an \mathcal{S} -reduction returns a nonzero polynomial f , we insert f into G . Initially, $G = \emptyset$, rather than a set containing $\{f_i\}$, since we do not know if f_i is \mathcal{S} -irreducible.

B is the set of pairs of the form (f, σ) , where f is a polynomial that we \mathcal{S} -reduce with respect to σ . Initially, we know that $\mathcal{S}(f_i) = e_i$; therefore, we initialize $B = \{(f_i, e_i)\}_{i=1}^m$.

The idea of the algorithm is to build an \mathcal{S} -Gröbner basis by finding its elements in ascending value of \mathcal{S} ; that is, always to choose $(f, \sigma) \in B$ such that σ is minimal. (See step 4c.)

Remark 20. In virtue of remark 19, for each σ we can keep in B at most one polynomial f such that $\mathcal{S}(f) = \sigma$. For the same reason we can, at any time, remove (f, σ) from B if we can find another polynomial f' such that $\mathcal{S}(f') = \sigma$ and $\text{LT}(f') < \text{LT}(f)$.

In practice we will remove from B a pair (f, σ) if we can find a $t \in \mathbb{T}^n$ and a $(f', \sigma') \in G \cup B$ such that $t\sigma' = \sigma$ and $t \text{LT}(f') < \text{LT}(f)$.

The pseudo code of the algorithm is the following:

Algorithm 21.

Input: $\mathcal{F} = (f_1, \dots, f_m)$: an element of P^m ,

μ : an ordering on \mathbb{T}^n ,

μ' : an ordering on \mathbb{T}_m^n .

Output: G : an \mathcal{S} -Gröbner basis of $I = (f_1, \dots, f_m)$.

(1) $L := \emptyset$

(2) $G := \emptyset$

(3) $B := \{(f_1, e_1), \dots, (f_m, e_m)\}$

(4) While $B \neq \emptyset$

(a) $B := \{(f, \sigma) \in B \mid \sigma \notin \langle L \rangle\}$

(b) Remove from B any (f, σ) such that we can find $(f', \sigma') \in G \cup B$, $t \in \mathbb{T}^n$ satisfying $t\sigma' = \sigma$ and $\text{LT}(tf') < \text{LT}(f)$

(c) Pick $(f, \sigma) \in B$ with minimal σ .

(d) $f := \mathcal{S}\text{-reduce}(f, \sigma, G)$

(e) If $f \neq 0$ then

(i) $B := \text{UpdatePairs}(L, G, B, (f, \sigma))$

(ii) $G := G \cup \{(f, \sigma)\}$

(f) Else

(i) $L := L \cup \{\sigma\}$

(5) Return $\{g : (g, \sigma) \in G\}$

Note that, since L may change during each iteration, some pairs we assumed to be normal turn out not to be normal. We remove those in step 4a.

In step 4b we implement the idea presented in Remark 20. Note that this is an optimization; the algorithm will successfully terminate without this line.

We still have to describe the two procedures Algorithm 21 invokes. The first is \mathcal{S} -reduce:

Algorithm 22 (\mathcal{S} -reduce).

Input: f : an element of I ,

σ : an element of \mathbb{T}_m^n ,

G : a set that contains the elements $(g, \mathcal{S}(g))$ of an \mathcal{S} -Gröbner basis with $\mathcal{S}(g) < \sigma$.

Output: f : an \mathcal{S} -irreducible polynomial with respect to σ .

(1) $f := f / \text{LC}(f)$

(2) While $\exists (g, \mathcal{S}(g)) \in G, t \in \mathbb{T}^n$ such that $t \text{LT}(g) = \text{LT}(f)$ and $t\mathcal{S}(g) < \sigma$

(a) $f := f - tg / \text{LC}(g)$

(b) If $f = 0$ then Return 0

(c) $f := f / \text{LC}(f)$

(3) Return f

This algorithm takes as input a polynomial f and a $\sigma \in \mathbb{T}_m^n$ and, as long as there is an \mathcal{S} -reducer for f in G , performs \mathcal{S} -reduction steps. Because of the hypothesis on G we know we obtain an \mathcal{S} -irreducible polynomial with respect to σ .

The second is *UpdatePairs*:

Algorithm 23 (*UpdatePairs*).

Input: L : a subset of $\text{LT}(\text{Syz } \mathcal{F})$,

G : a set that contains the elements $(g, \mathcal{S}(g))$ of a \mathcal{S} -Gröbner basis with $\mathcal{S}(g) < \sigma$,

B : a set that contains elements (g, σ_g) of polynomials that have yet to be considered,

$(f, \mathcal{S}(F))$: where $f \in I \setminus \{0\}$.

Output: B' : a set of pairs (f', σ) that satisfy Theorem 10, produced by the criterion.

(1) $B' := \emptyset$

(2) For each $(g, \mathcal{S}(g)) \in G$, if (f, g) is a normal pair

(a) Compute u_1, u_2 such that $\text{Spol}(f, g) = u_1 f + u_2 g$

(b) $\sigma := \max(u_1 \mathcal{S}(f), u_2 \mathcal{S}(g))$

(c) $B' := B' \cup \{(\text{Spol}(f, g), \sigma)\}$

(3) Return $B' \cup B$

Proposition 24. *Algorithm 21 terminates.*

Proof. First we show that step 4(f)i is executed only a finite number of times.

Because of step 4a, at a given time, we only consider σ that do not belong to L ; so when we execute step 4(f)i we really enlarge the P -module generated by L . Since P^m is noetherian this can happen only a finite number of times.

Also, that step 4(e)i is executed only a finite number of times. First note that if f is not primitive \mathcal{S} -irreducible (that is, f is only \mathcal{S} -irreducible), then Algorithm 23 does nothing, so no new polynomials are generated. In the proof of Proposition 14, we see that an \mathcal{S} -Gröbner basis contains only a finite number of primitive \mathcal{S} -irreducible polynomials. This completes the proof. \square

Theorem 25. *Algorithm 21 computes an \mathcal{S} -Gröbner basis of I .*

Proof. This is a direct consequence of criterion of Theorem 18: Previous remarks have shown that G contains only \mathcal{S} -irreducible polynomials, and the initial value of B ensures that the algorithm satisfies the first condition. For the second condition, for each normal pair (g_1, g_2) , we ensure that we have a polynomial f and a monomial t such that $\mathcal{S}(tf) = \mathcal{S}(\text{Spol}(g_1, g_2))$. \square

The fact that non-primitive \mathcal{S} -irreducible polynomials do not generate any new pairs plays a *central* role in this proof of termination. Without it, the thesis does not hold: if we drop condition (1) of Definition 16, it is possible that the algorithm could enter an infinite loop, computing an infinite number of polynomials of the form $t_i f$ where $\{t_i\}$ is an infinite set of terms and f is an \mathcal{S} -irreducible polynomial and each of the $t_i f$ is \mathcal{S} -irreducible itself. (This occurs, for example, in the implementation of [15].)

7. COMPARISON WITH PREVIOUS WORK

In this section, we consider how this algorithm is both similar and different to two algorithms in past work: the staggered linear basis algorithm of Gebauer and Möller [8] and the F5 algorithm of Faugère [5]. We also illustrate explicit differences on three particular examples.

7.1. Comparison with Staggered Linear Bases. The Staggered Linear Basis algorithm (in the rest of this section, SLB) [8] introduced a special kind of Gröbner basis.

Definition 26. The set $B \subset I$ is a *staggered linear basis* of the ideal I if for all $f \in P$

- if $f, g \in B$ and $\text{LT}(f) = \text{LT}(g)$, then $f = g$; and
- if $t \in \mathbb{T}$ and $tf \in B$, then $f \in B$.

A full review of SLB is beyond the scope of this paper, but it is worth comparing to the present algorithm because both use trivial syzygies to detect zero reductions. To facilitate the explanation, we temporarily adopt the notation $t_i = \text{LT}(f_i)$ and $t_{i,j} = \text{lcm}(t_i, t_j)$.

SLB tracks monomial ideals for each polynomial among the generators. Initially, we have

$$Z_i = (t_1, t_2, \dots, t_{i-1}).$$

Critical pairs (f_i, f_j) (with $i < j$) are rejected whenever $t_{ij}/t_j \in Z_i$. If instead the S -polynomial of (f_i, f_j) is computed, then Z_j is expanded by adding the ideal generated by t_{ij}/t_j . If reduction of the S -polynomial results in a new polynomial f_k being added to the basis, SLB also creates a new ideal

$$Z_k = (Z_j + (t_i)) : (t_{ij}/t_j) + (t_1, \dots, t_{k-1}).$$

Despite the use of principal syzygies in the initial definition of Z_i , a fundamental difference between the algorithms lies in the fact that SLB does not compute, let alone consider, the leading module term $\mathcal{S}(f)$ of any polynomials. So a polynomial can be \mathcal{S} -irreducible even if it is top-reducible, and the normal pairs of the F5 Criterion are not the same as the critical pairs of SLB. As a result, the approach in SLB behaves quite differently, and fails to detect certain zero reductions detected by F5 and the present algorithm.

7.2. Comparison with F5. At first glance, algorithm 21 may appear very different from the F5 algorithm. However, if we define μ' to be

$$te_i <_{\mu'} se_j \iff \begin{cases} i < j & \text{or} \\ i = j \text{ and } t < s \end{cases}$$

for any $t, s \in \mathbb{T}^n$ and $1 \leq i, j \leq m$, there is an interesting relationship between \mathcal{S} -Gröbner bases and $\text{LT}(\text{Syz } \mathcal{F})$. Define, for $1 \leq l \leq m$, $\pi_l : P^m \rightarrow P$ as the projection on the l -th component, then

$$(7.1) \quad \pi_l(\text{PSyz } \mathcal{F}) = (f_1, \dots, f_{l-1})$$

where $\text{PSyz}(\mathcal{F})$ is the P -module of principal syzygies, defined as $\text{PSyz}(\mathcal{F}) = \langle f_i e_j - f_j e_i \rangle_P \subseteq P^m$; $\text{PSyz}(\mathcal{F})$ is clearly a P -submodule of $\text{Syz}(\mathcal{F})$.

Suppose we have $f \in I \setminus \{0\}$ and we know $\mathcal{S}(f) = te_i$, for some $t \in \mathbb{T}^n$. It follows from the definition of \mathcal{S} that $f \in (f_1, \dots, f_i)$. Hence, (7.1) implies that

$$\text{LT}(f)e_{i+j} \in \text{LT}(\text{PSyz}(\mathcal{F})) \quad \text{for some } j \geq 1.$$

With this choice for μ' , we can improve the performance of Algorithm 21 by adding an instruction right after step 4(e)ii:

$$L := L \cup \{\text{LT}(f)e_{i+1}, \text{LT}(f)e_{i+2}, \dots, \text{LT}(f)e_m\},$$

where $\sigma = te_i$ for some $t \in \mathbb{T}^n$. In other words, whenever we find a new element of G , we also find new elements of L .

Also, due to the ordering on P^m , the structure of an \mathcal{S} -Gröbner basis G is very special. We find the elements of G in ascending value of \mathcal{S} : we first find all the elements g such that $\mathcal{S}(g) = te_1$ for some $t \in \mathbb{T}^n$, then those g such that $\mathcal{S}(g) = te_2$ for some $t \in \mathbb{T}^n$ and so on. Is easy to see that the real value of f_l is never considered in any computation, until the algorithm has finished producing all the elements of G with $\mathcal{S}(g) = te_i$ for some $t \in \mathbb{T}^n$ and $i < l$. If we make the further assumption that $\text{Syz}(\mathcal{F}) = \text{PSyz}(\mathcal{F})$, we conclude that the algorithm never reduces a polynomial to 0, since we discover every leading term of the syzygies in advance.

Therefore, we may say that, in this case, Algorithm 21 is *incremental*, as it first produces an \mathcal{S} -Gröbner basis of (f_1) , then an \mathcal{S} -Gröbner basis of (f_1, f_2) and so on, and avoids all the reductions to zero; this behavior is the same as Faugère's F5 algorithm.

We can couch the use of “simplification rules” in F5 [5, sect. 6], also called the *rewritable criterion*, in vocabulary similar to that used in this paper: F5's algorithm to compute S -polynomials (SPol) discards any (te_i, f) when

- there exists some other $(t'e_i, f') \in G \cup B \cup B'$ such that $t \mid t'$, and
- f' was computed *before* f .

This concept is related to Remark 19 in this paper; roughly we know we can “decide” how to obtain an \mathcal{S} -irreducible polynomial with a given signature. We prefer to start with a polynomial with the smallest leading term we know of, while in F5 just the first generated polynomial is kept.

This parallel carries over to the computation of L , which here is used to prevent the computation of any $te_i \in \text{LT}(\text{Syz } \mathcal{F})$ more than once. When a polynomial is reduced to zero in F5, the simplification rule is added even though the polynomial is discarded, and this rule ensures that any polynomial f with $\mathcal{S}(f) = t'e_i$, where $t \mid t'$, is not computed. In other

Algorithm	Number of zero reductions			
	MMT92	Cyclic-5	Cyclic-6	Katsura-5
Staggered Linear Basis	3	46	446	10
F5	0	0	16	0
Algorithm 21	0	0	8	0

TABLE 1. Number of zero reductions during execution of algorithms SLB, F5, and Algorithm 21.

Algorithm (size of red. GB)	Size of basis			
	MMT92	Cyclic-5	Cyclic-6	Katsura-5
Staggered Linear Basis	8	38	99	22
F5	10	39	202	30
Algorithm 21	10	39	155	30

TABLE 2. Size of the Gröbner basis computed by algorithms SLB, F5, and Algorithm 21.

words, F5 has an implicit provision for avoiding the computation of non-trivial syzygies, like the algorithm here.

7.3. Concrete examples. We examine how all three examples perform on three “standard” systems:

- the system “MMT92”, $F = \{yz^3 - x^2t^2, xz^2 - y^2t, x^2y - z^2t\}$ from [5] (this seems first to appear in non-homogenized form in [12]);
- the homogenized Cyclic-5 system; and
- the homogenized Katsura-5 system.

We consider

- (1) the number of zero reductions; and
- (2) the size of the Gröbner basis generated.

The tests were carried out in unoptimized implementations of each algorithm in Sage [16, 1, 14, 13], and are available online.

The results are in Tables 1 and 2. Table 1 shows that the Staggered Linear Basis algorithm computes some zero reductions *even though* the systems are regular sequences. Neither F5 nor Algorithm 21 computes any zero reductions except in Cyclic-6, which is not a regular sequence. In that system, Algorithm 21 computes a smaller basis, and it computes fewer zero reductions. This appears to be due to the fact that it proceeds by ascending signature (line 4c) rather than by ascending lcm (compare to algorithm Spol in [5]).

8. CONCLUSIONS AND FUTURE WORK

This paper has reformulated the F5 criterion, which in its original form is due to [5], and provided a new proof of this criterion’s correctness. We have introduced the ideas of \mathcal{S} -Gröbner basis and \mathcal{S} -irreducible polynomials, and have shown that if a set of polynomials G satisfies the F5 criterion, then G is an \mathcal{S} -Gröbner basis and not just a Gröbner basis. In this new setting, we were able to drop many restrictions present in [5]: we can freely choose any ordering on P^m , and there is no need to for the sequence (f_1, \dots, f_m) to be regular.

Our statement of the criterion is quite different from the original: we require that all the polynomials in the set G be \mathcal{S} -irreducible; we require that if $e_i \notin \text{NS}(\text{Syz } \mathcal{F})$, then there exist $g \in G$ such that $\mathcal{S}(g) = e_i$; and we impose a condition on the *signature* $\mathcal{S}(\text{Spol}(g_1, g_2))$, rather than the usual condition that

$$\text{Spol}(g_1, g_2) = \sum_{i=1}^{\#G} h_i g_i \quad \text{such that} \quad h_i \neq 0 \implies \text{LT}(h_i) \text{LT}(g_i) \leq \text{LT}(\text{Spol}(g_1, g_2)).$$

(Faugère calls this latter condition $o(\text{Spol}(g_1, g_2))$.) We also changed the definition of *normal pair* by adding a new condition: the fact that we can consider only primitive \mathcal{S} -irreducible polynomials.

We then proposed a simple algorithm to show an application of the new criterion. The algorithm presented here is mainly demonstrative, and does not include many “obvious” optimizations such as holding off on the computation of a new polynomial f until it is actually needed in step 4d of Algorithm 21.

The authors would like to thank the referees for helpful and instructive comments that improved the paper.

APPENDIX

In October 2011, Vasily Galkin of Moscow State University contacted us with a question about Proposition 14. His question was sparked by the definition of μ' as an *admissible* ordering on \mathbb{T}_m^n . Apparently, we used the wrong word; our reference for the notation (which, it amazes us to report now, we did not include in the bibliography) was [10]. (As far as we can tell, it is the only textbook that uses the word “monomodule”.) This text *does not define* an “admissible” ordering for a module; it defines either a module ordering (p. 54) or a compatible ordering (p. 55). It seems that when we wrote “admissible”, we meant “compatible”. Indeed, if the ordering is not compatible, the \mathcal{S} -Gröbner basis may be infinite, as the following example shows.

Example 27. Let $<$ be the degrevlex ordering with $x > y > z$. Let $<'$ be the module ordering with

$$xe_1 <' ye_1 <' xe_2 <' ye_2 <' ze_1 <' ze_2$$

extended to all other module terms in the following way: $te_i <' ue_j$ if

- $\deg t < \deg u$ (total degree), or
- $\deg t = \deg u$ and
 - $\deg_z t < \deg_z u$, or
 - $\deg_z t = \deg_z u$ and
 - * $i = 1$ and $j = 2$, or
 - * $i = j$ and $\deg_y t < \deg_y u$, or
 - * $i = j$, $\deg_y t = \deg_y u$, and $\deg_x t < \deg_x u$.

It is routine to verify that this is a module ordering; it is obviously not compatible with $<$.

Let

$$f_1 = x^2 + xy, \quad f_2 = xy + z^2.$$

We have

$$(8.1) \quad \text{LT}(\text{Syz}(\{f_1, f_2\})) = \{t \cdot \text{LT}(f_2 e_1 - f_1 e_2) : t \in \mathbb{T}^n\} = \{t \cdot z^2 e_1 : t \in \mathbb{T}^n\}.$$

For $i > 3$, let f_i be the S -polynomial of f_1 and f_{i-1} . We have

$$\begin{aligned} f_3 &= xy^2 - xz^2 \\ f_4 &= xy^3 + x^2z^2 \\ &\vdots \\ f_i &= xy^{i-1} + (-1)^i x^{i-2}z^2. \end{aligned}$$

For $i \geq 2$, $\mathcal{S}(f_i) = x^{i-2}e_2 = \min \{te_2 : t \in \mathbb{T}^n, \deg_t = i-2\}$. So f_i is not \mathcal{S} -reducible by f_2, \dots, f_{i-1} even though it is top-reducible by them. Additionally, $\text{LT}(f_1) \nmid \text{LT}(f_i)$, so f_1 is \mathcal{S} -irreducible.

It remains to see if the f_i are primitive. Let $g \in I \setminus \{0\}$ such that $\text{LT}(tg) = \text{LT}(f_i)$ for some $t \in \mathbb{T}^n \setminus \{1\}$. The factors of $\text{LT}(f_i)$ are x and y only, so $t = x^a y^b$ for some $a, b \in \mathbb{N}$. We also want $\mathcal{S}(tg) = \mathcal{S}(f_i)$; we claim that this is possible only if $t = x^a$ for some $a \in \mathbb{N}$. To see why, assume that $\mathcal{S}(tg) = \mathcal{S}(f)$. If $t\mathcal{S}(g) \neq \mathcal{S}(g)$, we must have $t\mathcal{S}(g) \in \text{LT}(\text{Syz}(\{f_1, f_2\}))$, which would imply that $t\mathcal{S}(g) = uz^2e_1$ for some $u \in \mathbb{T}^n$. Since $\mathcal{S}(g) \in \text{NS}(\text{Syz}(\{f_1, f_2\}))$, we infer that $2 > \deg_z \mathcal{S}(g) = \deg_z(t\mathcal{S}(g))$, a contradiction to $t\mathcal{S}(g) = uz^2e_1$. Thus, $t\mathcal{S}(g) \in \text{NS}(\text{Syz}(\{f_1, f_2\}))$, and $t\mathcal{S}(g) = \mathcal{S}(tg)$. So $t\mathcal{S}(g) = \mathcal{S}(tg) = \mathcal{S}(f_i) = x^{i-2}e_2$. As claimed, $t = x^a$ for some $a \in \mathbb{N}$. Since $t \neq 1$, $a > 0$; since $t \mid \text{LT}(f_i)$, $a < 2$. Hence, $t = x$ and $\text{LT}(g) = y^{i-1}$, but a computation of the Gröbner basis shows that $y^{i-1} \notin \text{LT}(I) = \langle y^2z^2, xz^2, x^2, xy \rangle$. Hence, there do not exist $g \in I \setminus \{0\}$ and $t \in \mathbb{T}^n \setminus \{1\}$ such that $\text{LT}(tg) = \text{LT}(f_i)$ and $\mathcal{S}(tg) = \mathcal{S}(f_i)$.

It follows that any \mathcal{S} -Gröbner of I with respect to $<$ and $<'$ must have all the polynomials f_2, f_3, \dots , and is thus infinite.

REFERENCES

- [1] Martin Albrecht and John Perry. Implementation of Faugère's F5 algorithm. Sage library, 2008.
- [2] M. Bardet. *Étude des systèmes algébriques surdéterminés. Applications aux codes correcteurs et à la cryptographie*. PhD thesis, LIP6, 2006.
- [3] B. Buchberger. *Ein Algorithmus zum Auffinden der Basiselemente des Restklassenringes nach einem nulldimensionalen Polynomideal*. PhD thesis, University of Innsbruck, 1965.
- [4] Bruno Buchberger. A criterion for detecting unnecessary reductions in the construction of Gröbner bases. In E. W. Ng, editor, *Proceedings of the EUROSAM 79 Symposium on Symbolic and Algebraic Manipulation, Marseille, June 26-28, 1979*, volume 72 of *Lecture Notes in Computer Science*, pages 3–21, Berlin - Heidelberg - New York, 1979. Springer.
- [5] J. C. Faugère. A new efficient algorithm for computing Gröbner bases without reduction to zero (F_5). In *ISSAC '02: Proceedings of the 2002 International Symposium on Symbolic and Algebraic Computation*, pages 75–83, New York, NY, USA, 2002. ACM Press.
- [6] Jean-Charles Faugère. A new efficient algorithm for computing Gröbner bases (F_4). *Journal of Pure and Applied Algebra*, 139(1–3):61–88, June 1999.
- [7] Justin Gash. *On Efficient Computation of Gröbner Bases*. Ph.D. dissertation, Indiana University, Bloomington, IN, 2008.
- [8] Rudiger Gebauer and Hans Möller. Buchberger's algorithm and staggered linear bases. In *Proceedings of SYMSAC 1986 (Waterloo/Ontario)*, pages 218–221. ACM Press, 1986.
- [9] Rudiger Gebauer and Hans Möller. On an installation of Buchberger's algorithm. *Journal of Symbolic Computation*, 6:275–286, 1988.
- [10] Martin Kreuzer and Lorenzo Robbiano. *Computational Commutative Algebra*, volume 1. Springer, 2000.
- [11] Daniel Lazard. Gröbner bases, Gaussian elimination, and resolution of systems of algebraic equations. In J. A. van Hulzen, editor, *EUROCAL '83, European Computer Algebra Conference*, volume 162, pages 146–156. Springer LNCS, 1983.

- [12] Hans Möller, Teo Mora, and Carlo Traverso. Gröbner bases computation using syzygies. In *ISSAC '92: Proceedings of the International Symposium on Symbolic and Algebraic Computation*, pages 320–328, New York, NY, USA, 1992. ACM.
- [13] John Perry. Implementation of Staggered Linear Basis algorithm. Sage library, 2008.
- [14] John Perry. Implementation of Arri's F5 variant. Sage library, 2010.
- [15] Till Stegers. Faugere's F5 algorithm revisited. Cryptology ePrint Archive, Report 2006/404, 2006. <http://eprint.iacr.org/>.
- [16] William Stein. *Sage: Open Source Mathematical Software (Version 4.1.1)*. The Sage Group, 2010. www.sagemath.org.

SCUOLA NORMALE SUPERIORE DI PISA - PIAZZA DEI CAVALIERI, 7 - 56126 PISA, ITALY

Current address: Microsoft Corporation

E-mail address: `arri@sns.it`

UNIVERSITY OF SOUTHERN MISSISSIPPI, HATTIESBURG, MS USA

E-mail address: `john.perry@usm.edu`